JOURNAL OF THE EXPERIMENTAL ANALYSIS OF BEHAVIOR 1998, **70,** 215–234 NUMBER 2 (SEPTEMBER)

GREATNESS AND MISERY IN THE TEACHING OF THE PSYCHOLOGY OF LEARNING

ARMANDO MACHADO AND FRANCISCO J. SILVA

INDIANA UNIVERSITY AND UNIVERSITY OF REDLANDS

Overshadowed by more popular disciplines, the study of learning seems to have lost its prominent place in the undergraduate psychology curriculum. In the first part of this essay, we argue that one reason for this state of affairs is the current content of psychology of learning courses, namely, its disproportionate emphasis on facts, procedures, and everyday examples at the expense of functional and conceptual investigations. In the second part of the essay, we outline an alternative approach to the teaching of learning, one that emphasizes basic *contents* such as the conceptualization of learning as a biological adaptation or the study of temporal regulation, critical *methodological issues* such as the logic of experimental designs or the difficulties of measuring behavior, and broad *epistemological problems* such as the role of hypothetical constructs, the advantages of quantitative reasoning, or the origins of knowledge and its integration. By using learning as a means towards more fundamental ends, the splendor of the discipline and its prominent place in the undergraduate curriculum may be restored.

Key words: learning, teaching, conceptual analysis, methodology, epistemology

Where Sancho sees windmills,
Don Quixote sees giants.
(A. Gedeão, Portuguese poet)

In a recent study, Bell and Goodie (1997) categorized 5,495 advertisements for positions that appeared in the *APS Observer Employment Bulletin* in the years 1991 through 1996. One interesting statistic that emerged from their analysis was that only seven (or 0.13%) of those advertisements were for positions in animal learning and behavior. This number increases only to 109 (or 1.98%) if we include all positions that sought to hire a specialist in the area generally categorized as learning and behavior (e.g., behavior analysis, experimental analysis of behavior, human learning, motivation, learning theory). In contrast, there were 553 (10.06%), 413 (7.52%), and 401 (7.30%) advertisements for jobs in social, cognitive, and developmental psychology, respectively. The preceding statistics suggest that the study of learning has fallen on hard times.

In the same vein, several colleagues from different areas of psychology, including learning, have commented that the psychology of learning is a stagnant, if not dying, subdivision of psychology, one that is confined to a few, narrowly defined domains (e.g., choice and the matching law), interested in arcane issues of little interest to the uninitiated (e.g., the distinction between occasion setters and discriminative stimuli, or the shaping of interresponse times), imprisoned in outmoded philosophies of science (e.g., logical positivism), isolated artificially by its jargon (e.g., excitatory-appetitive conditioning, or concurrent variable-interval variable-interval schedules of reinforcement), and, alas, one that has produced few significant findings in more than two decades. We also have heard from students that a course on learning is useful only to the extent that it acquaints them with tools and techniques to investigate more modern and meaningful phenomena. For example, the conditioning of the rabbit's nictitating membrane may be a convenient tool to understand the timing functions of the cerebellum, and the Morris water maze may be a convenient technique to unravel the role of the hippocampus in spatial memory.

There are many reasons for the preceding state of affairs. The "cognitive revolution," with its emphasis on information processing, rules, and representations, has shifted the interest from simple forms of behavior and learning, such as habituation and Pavlovian and operant conditioning, to complex phenomena such as memory, language, and con-

---

We thank Jim Craig, Susan Goldstein, Susan Jones, Lisa Peters, Bill Timberlake, and Ruhiyyih Yuille for their helpful comments and suggestions on earlier drafts of this article.

Please address correspondence to Armando Machado, Department of Psychology, Indiana University, Bloomington, Indiana 47405 (E-mail: amachado@indiana.edu) or Francisco Silva, Department of Psychology, University of Redlands, 1200 East Colton Avenue, P.O. Box 3080, Redlands, California 92373-0999 (E-mail: silva@uor.edu).

sciousness. The growth of the neural sciences in psychology departments and the general enthusiasm for the decade of the brain have shifted the interest from observable changes in behavior to its neurophysiological substrates. The rapid growth of the applied areas of psychology has continued to shift the interest from basic, animal-centered research to socially significant, human-centered investigations.

In this essay we focus on another reason for the current state of the field of learning, namely, its teaching. We believe that those of us who study learning too often have gone about the business of teaching the discipline in an isolated, unreflecting, stereotyped, and occasionally even perfunctory way, a way that, were it not also tragic, would be quite ironic, for the very community that investigates the phenomenon of learning has remained largely silent about its teaching. As one piece of evidence, consider that only about 20 articles (1.2%) even remotely related to the psychology of learning have been published in the *Teaching of Psychology*, a journal that has published more than 1,570 articles. Of these 20 articles, most are geared toward introductory psychology courses and consist of descriptions of classroom demonstrations and laboratories. Thus, if it is true that learning has had—and, like any other course, will always have—its share of great teachers (the late Fred Keller comes to mind) and that some well-known instructors have shared their thoughts about teaching in general (e.g., Michael, 1993), it is also true that there is a clear shortage of published studies about the teaching of learning per se. Perhaps, as Robinson (1979) stated, the psychology instructor "tends to judge the virtues of his subject to be self-evident, and considers requests for justification to be impertinent. Worse—and this is especially and lamentably so in psychology—he may not be clear in his own thinking about what these self-evident virtues are" (p. 4). As another piece of evidence, consider that our experience in four psychology departments has shown that students systematically rate textbooks of learning poorly, even those written by top researchers in the field. The most frequent remarks are that learning textbooks are dry, their content boring, and their language esoteric.

The preceding pieces of evidence also suggest that our lack of public reflection on the teaching of learning may have either helped to portray the discipline as an anachronistic subject or done nothing to change that portrait. In either case, it is clear that we have failed to continue to convince students of the importance of learning, to excite them about its investigation, and, more generally, to spawn a new generation of scholars interested in the subject. Not surprisingly, then, the number of jobs in the field has dwindled, the scientific progress in the area has slowed down, and the false perception that learning is only a convenient toolbox has been strengthened. The greatness of the discipline may be overshadowed by the unreflecting way we too often have been teaching it.

This is an unfortunate state of affairs because, perhaps unlike any other course in the undergraduate psychology curriculum, learning is ideally suited to address a broad range of important issues in the study of science in general and psychology in particular. Thus, in learning, students may explore content-based topics such as how to conceptualize learning as a biological adaptation, how to study the spatial and temporal controls of behavior, and how to investigate the nature of concepts. Students may also learn about or revisit methodological topics such as the logical underpinnings of different experimental designs or the difficulties of measuring psychological constructs. Epistemological topics such as the tolerance for, and the role of, hypothetical constructs, hypothesis and theories in scientific development, the distinctions between correlation and causation, the advantages of quantitative reasoning, the origins of knowledge and its integration, and the tension between basic and applied research goals, are also central to the subject matter of the psychology of learning. Rich in ideas and in explicit connections with many other scientific and philosophical areas of inquiry, learning may serve as a means toward the achievement of important goals in the undergraduate curriculum and as a means toward integrating and revisiting topics learned in other courses—thinking critically about behavior and cognition, time, space, and causality, scientific strategies and methods, the discovery of meaningful problems, and the invention and construction of their solutions.

The purpose of the present essay is to describe some of the problems in the way that learning is currently taught (Part 1) and to suggest potential solutions to these problems (Part 2). We focus on the content of the subject, not on the strategies, practices, and ‘‘tricks of the trade’’ followed by instructors who teach it. We analyze how major topics are justified, elaborated, and assessed in introductory courses and textbooks on learning, the types of skills these courses and textbooks seem to promote, and the issues and skills they either ignore or underemphasize. We conclude the essay with some recommendations on how to change the way learning is taught and thereby enhance its importance within the undergraduate curriculum.

To avoid misunderstandings, it is important to note that what follows is not a criticism of any particular *approach* to the study of learning, but rather of specific *aspects* of learning courses and *practices* of its instructors. Understandably, instructors differ in their talents, experiences, interests, and goals. They differ in how they conceptualize learning itself, in how they assess the state of research in the field, and in how they view the content of the learning course. Instructors also differ in the characteristics of the population of students they teach. Hence, some adopt a strong behavior-analytic approach (e.g., Pierce & Epling, 1995), others a cognitive-associative viewpoint (e.g., Klein, 1996), and still others present the course content more eclectically (e.g., Domjan, 1993). Some deal exclusively with animal behavior (e.g., Staddon & Ettinger, 1989), others with human behavior (e.g., Grant & Evans, 1994), and still others, in the spirit of Keller and Schoenfeld (1950/1995), use the former to try to understand the latter (e.g., Catania, 1998; Donahoe & Palmer, 1994). Some emphasize mostly the basic laboratory science (e.g., Mazur, 1998; Tarpy, 1997), whereas others emphasize exclusively applications and everyday examples (e.g., Baldwin & Baldwin, 1998). In itself, there is nothing wrong with different approaches and emphases to the teaching of learning. But, when a course suffers from the problems identified below and, as a consequence, (a) divorces terminology, procedures, and factual discoveries from the conceptual and theoretical issues that give them the significance they have; (b) emphasizes applications of learning principles and experimental findings to hot topics for the sole purpose of motivating the students, and without any concern for the fundamental issues that underlie these principles and findings; (c) presents each topic as a mini-science with its own technical vocabulary and interpretative principles without integrating the material and clarifying the logical geography of the different models, concepts, and experimental findings of learning; and (d) promotes the memorization of facts and definitions as ends rather than means and demotes a variety of other skills critical to scientific behavior, then the meaning of the phenomenon of learning and the broad significance of its study will be lost.

However, before proceeding, two explanations about the source and target of our concerns are necessary. First, the dearth of published studies and reflections on the teaching of learning has forced us to rely on personal experience more than it would be desirable. So that the reader can assess properly the representativeness of that experience, we state that our arguments were formed in the course of teaching undergraduates at Indiana University and the University of Redlands, reading textbooks written for this subject, examining their accompanying test banks, discussing the teaching of learning with colleagues at the aforementioned universities as well as with colleagues at the University of Manitoba, Columbia University, Johns Hopkins University, and Duke University, and examining copies of their tests. Second, although our primary goal is to address issues surrounding the teaching of learning, we believe that the problems identified below characterize many other undergraduate courses in psychology; therefore, the recommendations we propose for learning may have wide applicability. For this reason, we invite readers who are not primarily interested in the area of learning or its teaching to think about a corresponding set of critical issues that can be addressed with the contents of their courses.

PROBLEMS WITH THE CONTENT TAUGHT IN THE PSYCHOLOGY OF LEARNING

What instructors judge to be of utmost importance in a particular subject can be assessed by the questions they ask in quizzes, exams, and other means of evaluating students' knowledge (Slavin, 1997; Woolfolk, 1998). The content of the questions reflects what instructors believe to be the critical areas of research in their discipline, the main experimental findings, concepts, and theories, the significant problems that have been solved, and those problems for which more research or conceptual clarity is needed. If we adopt this viewpoint and then review test banks that accompany undergraduate learning texts such as Barker (1997), Klein (1996), Mazur (1998), Pierce and Epling (1995), to name a few, we are led to the reflections summarized below.

The Tyranny of Facts

Perhaps the most striking feature of these test banks is, in terms of content, how many questions assess only the most basic knowledge, the definition of words and procedures. In terms of form, it is striking how many questions convey the message, if not directly then by inference, that memorization of disconnected facts is the distinctive feature of scientific understanding. Consider, for example, the following test question:

> Descartes thought that the mind and the body interact through the
> (a) pineal gland
> (b) hypothalamus
> (c) thalamus
> (d) corpus callosum

What is the relevance of this question to a student of learning who has learned little about Descartes' general philosophy or his theory of the mind? Has the student learned anything substantive about the mind-body problem, Cartesian dualism, or its roots, supporters, and critics? What was the importance, even for Descartes, of the place where the mind interacts with the body? Has the student learned anything about the human nervous system, its various structures, and its multiple functions? Here is a second example:

> Under which of the following reinforcement schedules does the rate of responding increase just before the delivery of reinforcement?
> (a) fixed interval
> (b) variable interval
> (c) fixed ratio
> (d) variable ratio

What have students learned if they can identify the fixed-interval option in this question but do not know why reinforcement schedules are important, how changes in response rate are used to study the temporal regulation of behavior, and the like? If we agree that science depends on the sort of functional relation contained in this question, then we disagree that it is exhausted with them, or, more important, that it starts with them. Yet, test banks and study questions rarely go beyond the memorization of isolated pieces of facts and procedures.

Now, it is likely that not all instructors of learning place these sorts of questions on their exams, that some ask students to answer multiple-choice questions that require the coordination of functional relations and their application to new settings, and that others ask mainly essay-type questions. However, neither of these points allays our initial concerns because (a) if no instructors used these test banks, then publishers would stop paying authors to write these supplements to textbooks, (b) samples of exams from several learning courses revealed many questions of this type, and (c) essay-type questions do not guarantee that students will be required to go beyond a simple, definition-based view of learning. For example, neither of the questions below requires knowing more than basic definitions or procedures.

> Distinguish among positive reinforcement, negative reinforcement, punishment, and extinction.
>
> What is conditioned inhibition? Describe a procedure for producing a conditioned inhibitor.

We need not impugn the value of basic definitions and procedures in science to be troubled by the acceptance of the regurgitation of such information as the most important, and frequently the only, criterion of comprehension. The current emphasis on isolated facts at the cost of functional relationships and concepts, on tabular asterisks at the expense of theoretical risks as Meehl (1978) stated it, is also revealed by the fact that although many experimental procedures and results are described, few textbook authors attempt to go beyond the data and elaborate on the theoretical challenges posed by such results. For example, most textbooks describe the standard procedures of simultaneous, delay, trace, and temporal conditioning, but rarely do they ask the question: What conclusion, if any, can be drawn concerning the role of time in the acquisition of an effect due to the association between two events? Similarly, if most textbooks describe the well-known experiments on concept learning or categorization in pigeons, few discuss the critical issue of what exactly controls the learner's discriminated responding, the proposed alternatives, and their relative empirical or conceptual merits. As a third example, consider that if most textbooks describe the procedure of shaping by successive approximations, few discuss the conceptual and empirical problems raised by the multidimensionality of behavior that one tries to shape, the geography of the corresponding behavioral landscape, and the criteria to decide about the best trajectory to follow in a given case.

To illustrate what is critically missing from our undergraduate education, consider the following example. A student is using the method of successive approximations to condition a rat to stand on its hind legs and pull with its forepaws a chain suspended from the ceiling of an experimental chamber. The behavior of orienting the head towards the chain was selected as the first approximation to the target behavior, and after a few reinforcers the rat spent most of its time looking in the appropriate direction. Then, as the second step, the rat had to approach the area immediately underneath the chain to receive the reinforcer. Again, this goal was attained easily. In the next step, the student required the rat to rear and touch the chain with its forepaws. However, at this stage, progress was slowed considerably because the rat did not make a sufficient number of criterion responses. After repeated failure, the frustrated student asked a lab assistant for help in overcoming this obstacle. Digging into her repertoire of "tricks," the lab assistant recommended that the student try to shape rearing first, regardless of where it occurs in the chamber, and then proceed from there. Ten minutes after following her advice, the student had the rat pulling the chain reliably.[1]

Given this example, what sorts of *questions* could students ask? They could, for example, ask why it was necessary to shape rearing before shaping the animal to go to the location where rearing should occur. But how would they proceed from here? They could, for example, design a between-groups experiment and eventually attach a tabular asterisk to the group of rats that was shaped to rear anywhere in the chamber before they were shaped to rear near the location of the chain. But what would be the theoretical significance of the experiment? Perhaps, one may argue, this is a perfunctory experiment. Perhaps, we argue, the experiment should be preceded by an attempt to conceptualize what is involved in shaping (see also Sternberg, 1997).

One could imagine a three-dimensional behavioral space, a sort of topographical landscape contained in a cube (see Killeen, 1992). The initial behavior of the rat could correspond to one of the top corners of the cube and the target behavior to an opposite corner. The problem is to select the most effective path joining the initial and target locations. Presumably, not all paths require the same amount of travel time, nor are all paths equally easy to use (e.g., some may be "winding," others "as the crow flies"). There may be physiological structures, behavioral processes, or phylogenetic histories that are analogous to mountains and valleys, conditions that affect the relative ease of traversing different paths. What works best in one circumstance may be less effective in others. Using this metaphor, chain pulling is easier to reach via rearing than by some other path that does not first "pass through" rearing.

The point of the preceding example is not to suggest any specific theory or model for shaping nor to answer the puzzles raised by the phenomenon, but to illustrate a strategy to strike the chords of the students' imagination, to stimulate insightful conceptualizations and subsequent research. For example, how can we discover the geography of a behavioral landscape? What behavioral principles are revealed by the geography? How does it depend on the type of reinforcer used by the trainer? What implications, if any, does the study of shaping have for our teaching methods (e.g., the sequence of steps followed in training a complex skill)? Contrast this approach with the typical caricature-like presentation of shaping to undergraduate students—shaping is the scientific version of the hot-and-cold game played by children, or a method to train animals to perform tricks. No wonder that shaping is still referred to as more art than science (Midgley, Lea, & Kirby, 1989; Pear & Legris, 1987; Silva, in press).

---

[1] Students come already disposed to blame the rat for behaving incorrectly. A possible important lesson for them to learn is that the rat's behavior is always "right."

Insufficient Conceptual Analysis

Another problem with some contemporary textbooks and courses on learning concerns their unquestioned acceptance of loosely specified and poorly understood concepts and explanations. Unsatisfied with traditional approaches that explained behavior by relating current performance to past events, with theories that organized data, and with our current ignorance of behavior and brain function, many textbook authors have embraced wholeheartedly the words, concepts, and modes of explaining used by laypeople. For if evolutionary theory has legitimized some form of extrapolation from animal to human, the ramifications of the cognitive revolution also seem to have legitimized some form of extrapolation from human to animal. It followed that accounts of animal learning based on descriptions of human behavior have taken the lead in inspiring new experiments, in providing the logic to derive new predictions, and in interpreting empirical findings. The resurgence of this approach is somewhat puzzling given our hard-won battles against descriptions of animals enmeshed in a web of human activities, feelings, and controversies (e.g., Romanes, 1883; see also Blumberg & Wasserman, 1995; Timberlake & Silva, 1994). Yet, the content of some learning texts seems to suggest that learning theorists are now entirely clear about the meaning of beliefs, expectations, intentions, surprise, attention, memory, representations, thinking, and consciousness.

The passive diffusion of mentalistic terminology into the psychology of learning is due in great part to its rhetorical power. After approximately 20 years of using it as an explanatory device, beginning students of learning are likely to embrace these explanations without resistance (Dinsmoor, 1986). Moreover, students are likely to conclude that psychology is easy and largely a matter of common sense because the explanations they knew all along have now been stamped with the prestigious seal of scientific adequacy. Already overwhelmed by the amount of new empirical findings laid bare in front of them, students can now concentrate on what they perceive to be the ''hard stuff,'' the various experimental procedures and the practical significance of the data they have yielded. As Baldwin and Baldwin (1998) stated, ''Students are more familiar with everyday behavior than any other form of behavior; so it takes little effort to follow the examples, and they can focus their attention on learning the behavior principles and behavioral analysis'' (p. vi).

Lost in current practices, however, is the careful analysis of scientific concepts, their meaning, articulation, logic, grammar, and, alas, pitfalls (see the illuminating discussion of this topic in Kennedy, 1992, pp. 96–105). Wittgenstein (1958, p. 232) would certainly repeat his famous dictum:

> For in psychology there are experimental methods and *conceptual confusion*. . . . The existence of the experimental method makes us think we have the means of solving the problems which trouble us; though problem and method pass one another by.

Let us illustrate the lack of conceptual analysis in the psychology of learning with an example taken from experiments on blocking and the Rescorla-Wagner model. According to some characterizations of this model, stimuli have predictive values that change as a function of how surprised an animal is by the occurrence (or nonoccurrence) of an unconditional stimulus. It is argued that animals make predictions at the onset of a trial, compare them with the effective outcome of the trial, and then, on the basis of the degree of mismatch between predictions and outcome, revise their initial predictions. According to some performance rule, the predictive value of the stimuli determines observable behavior.

The concepts of *surprise, prediction,* and *value* refer to things that humans do and experience. For example, we predict sunny weather, with some degree of certainty, and are then surprised by the falling snow. But we may also be surprised by an event even if we made no predictions before it occurred. If someone who just left your house after an evening of socializing returns 10 minutes later, then you may be surprised by the event even if you made no prediction when you heard the knock at the door. In everyday language, then, surprise refers to a variety of states of affairs, to an event that occurred without being announced, for example, or to an event with a low probability given the current context. Typically, it does not refer to any mental state or operation of the nervous system. But the use of surprise in current theories of learning endows the concept with additional properties that are not warranted by its everyday accessions and that cannot be advanced without justification. According to some authors, surprise causes memory searches, refocuses attention, increases awareness of the environment, and motivates attempts to discover better predictors of the surprising event. That much is illustrated in the following account of blocking and unblocking, taken from an undergraduate learning text (Lieberman, 1993, pp. 458–459, italics added).

> When the noise was followed by shock during pretraining, the occurrence of the shock would at first have been *totally surprising*, leading to an *active search* for stimuli that could have predicted it. Once subjects learned that the noise was always followed by shock, however, the shock's occurrence would no longer have been *surprising* and thus would no longer have triggered a *memory search*. . . . When the light was added to the noise on the compound trials . . . at first subjects might not *have been certain* that the shock would still follow. On the first compound trial, therefore, the shock might have engendered some *degree of surprise* and thus a small amount of conditioning. Once *subjects realized* that the noise was still a reliable predictor, however, the occurrence of the shock no longer would have been *surprising*. With the possible exception of the first trial, therefore, no *memory search* would have occurred on any of the compound trials, and hence *no association* would have been formed between the light and shock. Thus, whereas an attentional analysis attributes blocking in this case to a *failure to attend* to the light, Kamin argued that *the rats were fully aware of* the light but that, because the shock was *not surprising*, the *rats made no attempt to form an association* between the two events.

What does it mean for a rat to be ''fully aware of,'' or ''certain that''? And who or what does the memory search: the rat, its mind, or its brain? Is the assertion ''Once subjects realized that the noise was still a reliable predictor . . . the occurrence of the shock no longer would have been surprising'' susceptible to empirical refutation, or is it true *by definition*? These questions are not meant to suggest that this account of blocking is incorrect, but to emphasize that we need to be clear about the grammar of the constructs used in our explanations. More generally, we may have yet to learn the lesson that worse than providing no explanation is providing the illusion of an explanation.

Motivation at the Expense of Understanding

We also note with concern some current textbooks' uncritical extensions of learning theory to everyday examples of human behavior (e.g., Baldwin & Baldwin, 1998; Grant & Evans, 1994). The reasons for these extensions are not hard to identify. Most instructors believe that students will enjoy and appreciate learning theory more if they see some of its potential to help society cope with significant problems (e.g., the treatment of phobias, or the education of the mentally disabled). Who would argue against the practice of extrapolating from laboratory studies to real-life examples and applications to facilitate understanding, or against the underlying assumption that, all other things being equal, motivated students learn more than unmotivated ones?

The problem, however, is that not all things remain equal. The presentation of everyday examples of a phenomenon often seems to supersede a careful analysis of a scientific issue, of its importance and difficulties. In the process, the study of functional relations, of how variables are constrained, of how hypotheses develop and theories are built is frequently overlooked, if for no other reason than time constraints. Thus, in one textbook the significance of the concept of superstition is discussed without any reference to the important issues raised by the work of Skinner (1948), Staddon and Simmelhag (1971), and Timberlake and Lucas (1985) except in the following cryptic endnote: ''Superstitious learning was the first effect of noncontingent reinforcement to be studied (Skinner, 1948, 1957), but other work shows it is much more complex than was originally thought (Hammond, 1980; Killeen, 1978; Schwartz, 1984, pp. 154–173; Staddon, 1977)'' (Baldwin & Baldwin, 1998, p. 312). When the difficulties of a topic are eliminated (by the instructor or textbook author, not the student); when a simplified picture is drawn (by the instructor or textbook author, not the student); and when relevant questions are posed with potential answers readily available (again by the instructor or testbank author, not the student), then it is clear that not everything has remained equal.

Consider another example, the topic of reinforcement schedules. Sometimes the topic is introduced as a simple extension of continuous reinforcement: ''The most basic schedule involves providing continuous reinforcement for each trial or response. A variety of other types of schedules provide partial or intermittent reinforcement'' (Drickamer, Vessey, & Meikle, 1996, p. 191). Typically, the textbook then illustrates the basic schedules with common, everyday life examples. Ratio schedules instantiate contingencies similar to those present in gambling devices or in piece work, and interval schedules to those involved in checking if the mail has arrived, watching for water to boil, or waiting for a city bus. In other cases, the human examples not only illustrate but also justify the study of reinforcement schedules. But what is less often stressed in textbooks (and probably by instructors as well) is that schedules of reinforcement are important for reasons that have little to do with everyday occurrences of the type illustrated by the preceding examples. First, they invite researchers to chart systematically the vast territory of the variables that control behavior, their interrelation, and their quantitative effects (Sidman, 1960). Time between reinforcers and response number, their average and their variance, are examples of the many variables that need to be considered to understand schedule performance. Second, reinforcement schedules bring to the forefront the problems of how temporally extended response patterns develop, how different responses compete, how the integration of events over time affect behavior, and how different actions are organized. How do we explain the acquisition of a fixed-interval scallop or the scalar property of temporal regulation observed under this and similar schedules? Stating the broad goals of schedule research, presenting the various problems raised by the study of schedule performance, and advancing some solutions to these problems (or even better, asking the students to propose their own and discussing them critically afterwards) may promote more understanding of reinforcement schedules than any exhaustive list of schedule names and their stereotypical effects, or any premature jump to everyday applications.

Our message should be clear now. If reinforcement schedules are sufficiently important to be included in a learning textbook or an undergraduate course, then we cannot justify their study mainly on the basis of simplistic arguments and loose analogies. Rather, we need to explain how reinforcement schedules lead us to investigate the complexity of many environmental contingencies and their effects on behavior. The same message applies to any other topic of study. More generally, we are not objecting to the use of animal experiments as a means of understanding human behavior, nor are we objecting to a course or textbook whose major goal is to understand everyday human phenomena (e.g., Donahoe & Palmer, 1994). In fact, we admire Keller and Schoenfeld's (1950/1995) classic book even though the authors explicitly avoided discussing facts that shed ''no light on human conduct'' (p. 2). What we are objecting to are mechanical, unexamined, and hence superficial transpositions of concepts from the animal laboratory to everyday contexts. When such transpositions are taken for granted (i.e., when they cease to be a problem), then the students' motivation may well increase but only in inverse proportion to their understanding of the relevant issues.

Lack of Integrative Reasoning

Contemporary practices in the undergraduate teaching of learning, perhaps in the teaching of all psychology, emphasize two basic skills, reading the textbook in search of definitions and then using them to select the correct answer from a pool of alternatives. What is conspicuously absent are the skills involved in commenting on a solution to a problem and generating alternative solutions, abstracting the key features of a definition and identifying violations of the grammar of a concept, judging the conceptual soundness of a hypothesis or theory, drawing a graph to depict a functional relation and interpreting it afterwards, applying a concept or theory in a new situation, wrestling with a problem for a few days and learning to consider it from different perspectives, structuring and writing an essay.[2] All these activities are essential to promote scientific reasoning, and yet it seems to us (and several colleagues) that they are not encouraged as much as they should be in typical undergraduate teaching. In fact, many of our students report great difficulties when asked to perform these various activities—''I'm lost'' is a common remark.

In what follows, we illustrate some of these difficulties with an example, and then describe how we have attempted to deal with them. Let us begin by assuming that a student responded correctly to the following (or similar) multiple-choice question:

> In which of the contingencies below does the instrumental response produce an appetitive stimulus?
> (a) punishment
> (b) negative reinforcement
> (c) omission training
> (d) positive reinforcement

Now ask yourself if that student would be prepared to solve the following conceptual ''puzzle'':

> Psychologists define positive reinforcement as follows: (a) a response produces a stimulus (i.e., the stimulus is a consequence of the response), (b) the response occurs more often, and (c) the response occurs more often because of the response–consequence relationship, not for some other reason. Although nonpsychologists are generally aware of the first two points, they often overlook the third. To see its relevance, consider the following example.
>
> A little boy is engaging in a temper tantrum that consists of, among other responses, crying and yelling. In an attempt to stop the tantrum, the father spanks the child. However, the result is that the child cries and yells even more, which leads to more spanking. A nonpsychologist reasons that spanking was a consequence of crying and yelling, so point (a) of the definition described above is satisfied. The responses, crying and yelling, occurred more often, so point (b) is also satisfied. Hence, spanking functioned as a positive reinforcer for crying and yelling.
>
> Given the above definition and scenario, answer the following questions: What, if anything, is wrong with the nonpsychologist's reasoning? Also, design an experiment that would show whether a stimulus is a positive reinforcer. That is, what conditions are necessary to show that a stimulus is a positive reinforcer?

Repeatedly, our experiences have shown that the verbal knowledge and skills required to answer the preceding multiple-choice question do not help students to solve the puzzle. To accomplish this latter task, students need to distinguish first the eliciting function of stimuli from their reinforcing function. Then, and only then, can they design an experiment to determine which function is present in the example. Then, and only then, can they relate this distinction to others introduced in the course (e.g., closed- vs. open-loop control systems). Let us illustrate the sort of answer we expect from our students.

> The nonpsychologist's reasoning is incorrect because, in all likelihood, the child is crying more vigorously because the spanking itself elicits, not reinforces, crying. In other words, that spanking is contingent on crying is irrelevant to explain the crying. An experiment to test the preceding interpretation might consist of the following. Imagine a thought experiment in which two groups of children are treated in the following cruel way: Members of one group are spanked when they begin to cry (experimental group); members of another group are spanked with similar frequency but regardless of their behavior (control group). If we observe that the children in both groups cry, then we may conclude that crying is not due to the reinforcing effects of spanking but to its eliciting effects.
>
> In general, to show that an event is a reinforcer, we need to show that the response increases in frequency because it is followed by the event, not because of any other property. Suppose then that you are trying to evaluate whether Event E is a reinforcer of Response B. To do so, we would need two groups. Group

---

[2] We are often surprised by how impatient students become when they cannot solve a problem in the first 20 minutes or so of studying. Cantril (1977, p. 182) remarks that when Pavlov was asked by his students how they could be as creative as he was, he replied ''Get up in the morning with your problem before you. Breakfast with it. Go to the laboratory with it. Eat your lunch with it. Keep it before you after dinner. Go to bed with it in your mind. Dream about it.'' The ability to persevere in the face of difficulty is another critical but underemphasized skill.

> I would be presented with Event E whenever the participants engaged in Behavior B; Group II would be presented with Event E regardless of their behavior. If the participants in Group I display Behavior B more frequently than those in Group II, then we can conclude that Event E is a reinforcer. However, if the two groups display Behavior A equally, then Event E is not a reinforcer of Response B.[3]

This is a more difficult but also more instructive question than most multiple-choice or simple definitional questions. This is not to say that learning definitions is unimportant. Obviously, definitions and terminology form one of the basic building blocks of any science, but, in the same way that knowing some French words is not the same as speaking, writing, and reading French, to regurgitate definitions and facts is not the same as to understand a subject. In this respect, our position is similar to that of Menaechmus, the tutor of Alexander the Great. When asked for a shortcut in the study of geometry, Menaechmus replied, ''Oh King, to travel in your country there are roads for the royal family and roads for the common citizen. But in Geometry, there is only one road.'' The same is true of most subjects that are worth learning.

THE PSYCHOLOGY OF LEARNING: A MEANS TOWARD IMPORTANT ENDS

In the preceding section we discussed what we perceive to be some shortcomings in the teaching of the psychology of learning, though it is likely that they also apply to other psychology courses. Our views were formed in the course of teaching undergraduates, reading textbooks, discussing the teaching of learning with colleagues, and examining test banks. From these experiences we conclude that we are not teaching learning in a manner that highlights its numerous positive attributes, in particular the broad significance of the issues it raises, not only to students of psychology but to those of other sciences. The importance of using the course content toward broader and more important ends cannot be overemphasized, for the person who takes away from a ''course nothing more than a large body of disconnected and sketchily examined items of fact, method, or theory has only a superficial and temporary advantage over the person who never attended the course'' (Keller & Schoenfeld, 1950/1995, p. lxiii). Hence, our goal in this section of the essay is to discuss some of the broader and important issues that learning is well suited to illustrate.

Like any other domain of psychology, learning is far from a coherent, systematic, and consensual (or even consensible) body of knowledge. However, despite these features, the learning course is an excellent vehicle—perhaps better than most other courses—for teaching students to think like scientists, because learning has had a long history of factual, functional, and conceptual investigations. Within it have dwelled some of the deepest philosophical controversies of the century and the most important, reliable, and orderly findings of psychology, many of which have helped us to cope with challenging practical problems. The history of the study of learning has endowed it with a set of rich overarching themes. We propose to bring at least some of these themes to the forefront of the learning course.

We begin by discussing content-based issues, such as the explanation of biological adaptation and the central role of time in behavior. We then discuss methodological issues, such as the logic of experimental designs and the difficulties of measuring behavior. We conclude with a discussion of epistemological issues, such as the tolerance for hypothetical constructs, the differences between correlation and causation, and the advantages of quantitative reasoning. Although our set of examples is to some extent arbitrary—instructors with other purposes, interests, and expertise, or who are teaching different populations of students, are likely to select different ones—they serve as a vehicle to analyze some critical thematic issues that are not arbitrary.

Content-Based Issues

The study of learning is part of the study of a central problem of biology: adaptation. In addition to morphological and physiological adaptations (e.g., the structure of the human hand, the immune response to an antigen), all organisms show behavioral adaptations. Some of these adaptations are quite specific (e.g., the intricacies of the bee dance), but others are quite general (e.g., habituation). What is important to note is that adaptation is a problem to be explained, not a concept to be invoked only as an explanation (Pierce & Epling, 1988). To that extent, it may be useful to divide adaptations into ultimate or phylogenetic, and proximal or ontogenetic. The former come about through changes in the frequency of certain traits in a population as a result of natural selection; the latter result from changes in the frequency of certain kinds of behavior of an organism due to actions in, and interactions with, its environment. But the behavioral changes are themselves regulated by processes with evolutionary histories. Hence, learning is the study of a major product of phylogenetic adaptation and of a major process of ontogenetic adaptation.

---

[3] The student can then check his or her understanding with a new example: Suppose that, while I am salivating, you give me a piece of dry food. You see that I now salivate more and therefore you give me more dry food; I salivate even more. Can you conclude that dry food reinforced my salivation? Explain.

Evolution, learning, and adaptation. The study of learning and evolution may proceed along at least two paths, the first emphasizing content and the second form. Concerning the former, the phenomenon of so-called misbehavior, for example, illustrates clearly the boundary where evolutionarily selected adaptations limit, and sometimes compete and interfere with, learned adaptations (Breland & Breland, 1961). Thus, a raccoon trained to place a coin-like token in a piggybank for food reinforcement learns the task initially, but then begins to rub the tokens together rather than drop them in the bank. A chicken trained to step up to homeplate and hit a miniature baseball eventually starts to chase its own hits into the outfield rather than run around the bases.

The converse case occurs when learning processes sustain forms of behavior that have no obvious adaptive function or even seem to incur a net biological cost to the organism. For example, people sometimes engage in drug abuse, use contraceptives, care for adopted children, or commit suicide, all of which are activities that are difficult to understand from an evolutionary analysis because they seem to lower fitness. But, if we realize that learning processes increase the fit between *local* environments and behavior and that short-term success often correlates with long-term reproductive success, then we can see how these seemingly maladaptive activities may be maintained (for a somewhat different approach to the same issue see, e.g., Pinker, 1997).

The study of learning may also profit from the analysis of the logic or form of evolutionary theory. Constructs such as selection pressure, types of selection, units of selection, variation, and differential reproduction (see, e.g., Dawkins, 1976, 1986; Falconer & Mackay, 1989; Roughgarden, 1996) may help us understand the dynamics of learning (e.g., Campbell, 1960; Donahoe & Palmer, 1994; Plotkin, 1993; Popper, 1972; Skinner, 1981; Staddon & Simmelhag, 1971).[4] If we view different responses as analogous to a population of individuals and reinforcement and extinction of these responses as analogous to a fitness function defined over the individuals (cf. Glenn, 1991), then the processes of shaping and response differentiation, for example, may be considered to be a microcosm of evolution by means of natural selection. In a similar vein, some of the problems in evolutionary biology concerning the origins and the maintenance of inheritable variations have striking formal parallels to problems in learning theory concerning the origins and maintenance of behavioral variability. In a burgeoning domain of research, several investigators have shown that reinforced responding is not inevitably stereotyped; instead, variable responding can result from the reinforcement of infrequent responses (frequency-dependent selection) or variations in responding (e.g., Machado, 1992; Page & Neuringer, 1985; Schwartz, 1982).

Time and temporal regulation. For centuries, philosophers have argued that an analysis of time and space is essential to an understanding of causality and its related concepts of association, prediction, contiguity, and contingency (see Cook & Campbell, 1979). Not surprisingly, then, time is also essential in the study of learning for two complementary reasons. First, the very definition of learning points to changes in behavior across time (in contrast, the importance of time in other areas of psychology is often reduced to before-and-after ''snapshots'').[5] Second, learning also includes the study of the timing abilities of organisms. Thus, the undergraduate learning course is an ideal means for illustrating the general importance of time in the analysis of behavior.

---

[4] Many other authors have expressed similar views. For example, "And, for at least a century, it has been recognized that even the highest-known biological function, human thought, involves random generation of many alternatives and is only shaped up into something of quality by a series of selections. Like the elegant eyes and ears produced by biological randomness, the Darwin Machine's final product (whether sentence or scenario, algorithm or allegory) no longer appears random because of many millisecond-long generations of selection shaping up alternative sequences off-line" (Calvin, 1987, p. 34).

Pavlovian and operant conditioning, with their emphasis on acquisition and asymptotic responding and on transient and steady states, place the importance of the temporal dynamics of behavior front and center (e.g., Cooper, 1991). Thus, for example, initial presentations of an arbitrary stimulus (e.g., a tone) followed by a biologically relevant stimulus (e.g., food) do not seem to affect the behavior of a hungry dog in a meaningful way; however, after several temporally spaced pairings of these stimuli, the dog shows food-anticipatory responses during the previously neutral stimulus. The meaning of the tone, its ability to control behavior, has changed across time as a result of a multiplicity of temporal contiguities that, from the experimenter's viewpoint, constitute a contingency. The role of time is also critical to understanding choice, whether it is the initial sampling phase of the process or its terminal equilibrium phase.

The study of timing itself, of how animals and humans adapt to periodicities, durations, or temporal intervals that signal important events, constitutes an exciting domain of contemporary research in learning (Church & Meck, 1982; Gibbon, 1977; Killeen & Fetterman, 1988; Machado, 1997). By analyzing theories and experiments on timing, students learn to understand the basic processes of temporal control and how these processes contribute to more complex forms of behavior (e.g., self-control). They may also learn to evaluate the pros and cons of using constructs such as internal clocks, behavioral states, pulses, pacemakers, accumulators, comparators, and time-gates to account for behavior.

Time and the nature of historical accounts of behavior. Some of the obstacles confronting the beginning student of learning include the technical language of the discipline, its multiple taxonomies, and its artificial procedures (Dinsmoor, 1989). Although some of these obstacles also exist in other courses (e.g., students of geometry have to learn the meaning of bisector angles, alternate internal and external angles, similarity, isometry, etc.), the problem is compounded in the study of learning because the discipline's technical language, taxonomies, and procedures are rarely used in other courses. Hence, without proper justification, these roadblocks will not be overcome easily. Moreover, the very logic of explanations and interpretations based on learning principles needs to be made explicit and directly taught, because, as we mentioned above, the typical undergraduate student brings with him or her a radically different set of concepts and approaches to account for human and animal behavior. Here then is another opportunity to emphasize the broad issues raised in the study of learning.

Most learning theorists believe that an animal's behavior is a complex function of current context, the stimulus situation facing the animal, and its history, the trajectory the animal followed in the space of possible behavioral histories. In other words, behavior is understood not necessarily through internal structures, whether neurophysiological or cognitive, but through the current context and the sum of the effects of past events (reinforcers, punishers, verbal instructions, etc.). Current context and past history are thus the elementary dimensions of behavioral accounts based on learning theory. But the role played by current context is itself a function of the animal's history and, therefore, ultimately history is the most primitive dimension (see also Donahoe & Palmer, 1994).

An historical explanation requires knowledge of the various forms of environment–behavior interactions because, from an analytical viewpoint, the animal's history is simply the temporal integration of its interactions with the environment. But if one does not know the nature of what is integrated, and thus of what constitutes history, how can one develop an historical account? In other words, if one has not described and classified first, how does one know what needs to be explained? This basic issue may justify why psychologists who study learning have devoted considerable research to identifying and characterizing elementary types of environment–behavior interactions. In fact, the current taxonomies of forms of learning (e.g., classical and operant conditioning) and the functions of stimuli (e.g., conditional, reinforcing, discriminative) may be conceived as first steps toward the identification of the minimal units of a behavioral explanation.

---

[5] The exception is obviously developmental psychology, though the time epochs in this area are usually longer than those studied in learning.

Concerning the logic of the accounts based on learning principles, they are analogous to the method followed by Darwin to explain the formation of vegetable mold through the action of earthworms. When a certain Mr. Fish remarked that the weakness and size of such creatures were incompatible with the enormous task Darwin had assigned them, Darwin replied as follows: "Here we have an instance of that inability to sum up *the effects of a continually recurrent cause*, which has often retarded the progress of science, as formerly in the case of geology, and more recently in that of the principle of evolution" (Darwin, 1881, p. 6, italics added).

At this light, then, one can understand why, for learning theorists, time is a necessary condition for the emergence and the accumulation of the effects of small causes. Moreover, one can also understand why, following Bacon, many learning researchers have stressed the direct, often laboratory-based, reconstruction of behavior as a means of testing their knowledge of the factors that may control it (Smith, 1992).

These three sets of examples suffice to illustrate the importance and wide generality of the content-based topics studied in learning. We can bring to light the importance of these topics by unburying them from the piles of procedures, experiments, and techniques that currently hide them. If we then confront the student's conceptions of these issues with empirical findings and relate them to other approaches such as evolutionary biology and philosophy, then we will have made significant progress in using the content of the psychology of learning as a means toward these broader ends.

Methodological Issues

Ideally, in any psychology course, students should learn something about the game of science. To this end, the psychology of learning is the perfect course to learn about variables (the minimal pieces of any scientific explanation), functional relations (how these pieces constrain each other), and theories (the rules that describe how the pieces combine to form complex configurations). Students also learn the social rules of the game, how it is grounded on systematic replication, on parsimony, and on rigor and clarity in thought and expression. These rules provide the primary justification for the importance of methodological issues in science. In what follows, we identify some of these methodological issues, ones that in our view should be highlighted in learning because they are critical to the broader goal of understanding the game of science. The examples also show how methodological issues can be studied in their original (and natural) setting; that is, within specific content-based issues.

Apparatus and the measurement of behavior. Although many areas of psychology rely heavily on pencil-and-paper tests (or their microchip equivalents), human or computer data-collection systems, or self-report and interview methods, researchers in animal learning have continually been involved with the invention, design, and refinement of new apparatus and ways to measure behavior (Martin & Bateson, 1993; Timberlake & Silva, 1994). Within the study of learning, the never-ending search for a better mousetrap is based on a fundamental, albeit generally implicit, understanding of the organism being studied and on the information the researcher is explicitly seeking to obtain. In other words, the development of an apparatus shows the interplay, and sometimes the tension, between contingency-shaped behavior on one side and rule-governed behavior on the other side (Skinner, 1969). For example, the development of the maze was largely due to Small's (1899) observation that a burrow-like system would be ideal for studying rodent behavior. However, Skinner (1956) noted that rats exhibited too many extraneous activities in mazes, which interfered with his attempts to study eating behavior. To develop a reliable index of feeding, one that was easy to record and that could be modified to produce vigorous and orderly strings of learned behavior, Skinner gradually developed the operant chamber, used a new measure of feeding behavior (the rate of lever pressing), and invented a new recording device (the cumulative recorder). The development of the radial-arm and water mazes arose out of an interest in studying spatial memory processes in animals during appetitive and aversive circumstances (Morris, 1981; Olton & Samuelson, 1976), processes that are harder to examine in the typical operant chamber.

Perhaps most important, however, is that all refinements in the measurement of behavior, a process that Timberlake (1990) has called *tuning*, are related to our knowledge of the evolved characteristics of the species in question. The tuning process itself is a cornucopia of information about an animal's natural behavior and tendencies (Timberlake, 1990). Thus, a lever placed too high above the chamber floor evokes nosing the lever upward rather than pressing downward. This is not an unfortunate by-product of bad design, but instead illustrates that the lever press is tapping into a behavioral repertoire of manipulation responses, whereas a lever nosed upward may be tapping into a repertoire of investigatory responses (Timberlake & Silva, 1994; see also Stokes & Balsam, 1991).

Thus, by studying the reasons for the development of an apparatus, its strengths and limitations, in addition to the sorts of data it provides, students could be learning two important lessons concerning the game of science. First, as Bachelard (1975) once remarked, microscopes—and, we might add, operant chambers and radial-arm mazes—were invented to extend reason, not vision. Second, scientists are also epistemologists whenever they observe and think about the act of scientific observation. We are reminded of Sir Arthur Eddington's (1939) words:

> Who will observe the observers? The answer is—the epistemologist. He watches them to see what they really observe, which is often quite different from what they say they observe. He examines their procedure and the essential limitations of the equipment they bring to their task, and by so doing becomes aware beforehand of limitations to which the results they obtain will conform. (p. 21)

Experimental designs. In the study of learning, research methodology is always in the forefront (Timberlake & Silva, 1994; Warner & Warden, 1927). The study of Pavlovian conditioning relies heavily on large *N* between-groups designs, whereas the operant tradition relies strongly on small *N* within-subject designs (Kazdin, 1982; Sidman, 1960; see Bernard, 1865/1927, for an early application of this approach in the medical sciences). These are not just distinct research methods, but different approaches to the problems of the replication, reliability, validity, and irreversibility of observations (Cronbach, 1957). The two methodologies subsume different conceptualizations of the role of chance in scientific observations, for example, and assign different weights to the statistical and the experimental forms of control of behavioral variability. For instance, according to the small *N* tradition, variability within and among individuals is considered primary data to be used in the processes of improving behavioral control techniques and of searching for functional relationships between independent and dependent variables. In contrast, the large *N* tradition views variability as "unexplained variance," the "noise" inherent in nature. Similarly, the small *N* tradition views sequence effects and the irreversibility of behavior not as a problem to be overcome by experimental design (as the large *N* tradition does), but an effect to be studied and understood directly. The student who takes a learning course is exposed to these two methodological traditions and is left to organize and evaluate the merits of each (e.g., What sorts of issues recommend the use of between-subjects designs? When are functional relations obtained from group data representative of individual performance?) and to appreciate the importance of the intensive study of the individual (the idiographic approach) to our understanding groups (the nomothetic approach) (Allport, 1961).

The preceding examples illustrate how the psychology of learning raises general, important, and interesting methodological issues. If we succeed in highlighting them, even if at the cost of omitting an observation, procedure, or experiment, and invite our students to think about the relationships among apparatus, research methods, observation, the evolved tendencies of animals, and the explicit and implicit (via tuning) analyses of behavior, then we will again have used learning to provide an understanding of broader issues.

Epistemological Issues

Although the potential of the psychology of learning as a means of understanding epistemological issues has already surfaced in a variety of contexts (e.g., when we discussed the need to evaluate the foundations of scientific observations or the need to examine the sense or intelligibility of scientific concepts and statements), in this section we describe three additional examples. The first concerns the distinction between correlation and causation and is based on the view that behavioral adaptations embody knowledge of the correlational and causal textures of the world. The second concerns the importance of quantitative reasoning in science, and is based on the view that inventing preferably quantitative hypotheses about how variables are constrained is the first creative step in a scientific investigation. The third concerns the role of hypothetical constructs in scientific explanations, and is based on the view that students of science have to eventually reflect on the status of unobservable phenomena in science and on the rules to move from them to observable phenomena and back.

Correlation and causation. Through sensitivity to covariations or correlations, organisms adapt to the regularities in their environment. The study of Pavlovian learning is very much the study of the behavioral effects of correlations between events. However, correlations are not causations because the latter require control over the presumed cause of the event. In a sense, the study of operant conditioning is the study of causation because the organism is left with one degree of freedom: its own behavior. By varying its acts and experiencing their consequences, the organism generally adapts to the causal texture of its world. The distinction between Pavlovian and operant conditioning may thus be conceived as a distinction of control: when organisms have it versus when they do not.[6]

More generally, by relating operant conditioning to the study of causation and Pavlovian conditioning to the study of correlation, we can then begin to discuss issues of broad scientific importance; for example, the necessary and sufficient conditions for determining a causal relationship. For instance, imagine two situations, one in which there is a 5-s delay between a low-probability response and a reinforcer and another in which there is a similar delay but the response has a high-probability of occurring. Will the organism's responding be affected by the response–reinforcer causal relation in both cases, and to the same degree? How does it happen that the required response, not the behavior that fills the 5-s gap, comes to be increased by the reinforcer? What do our answers to these questions reveal about operant conditioning in particular and the conditions for ascertaining cause–effect relationships in general?

Quantitative reasoning. Quantitative concepts and the ability to reason with them are trademarks of science and scientific maturity (e.g., Carnap, 1966). We are not referring to the mathematical skills needed for tasks such as the derivation and solution of complex equations, but to the understanding of the most primitive concepts of science—the concepts of variable and functional relations (Whitehead, 1948). Without a working knowledge of these concepts, no further progress can be achieved in mastering the game of science.

And yet, a significant number of students of the psychology of learning, and, we suspect, of psychology in general, do not know how to use appropriately the concepts of variable and functional relations. If these were learned in other courses, then they have been forgotten. Repeatedly, in our experience teaching learning, we encounter students who can state verbally the correct general trend of a functional relation (e.g., the rate at which a hungry rat presses a lever decreases as the delay between the response and the delivery of a pellet of food increases) but do not know how to draw that relation in a graph. The converse difficulty also occurs; after drawing a curve, a student does not know how to interpret it.[7] These difficulties are generally accompanied by problems understanding the concept of *variable.* For example, many students of learning are not aware of the often natural limits of variables (i.e., that there is only so much that a dog can salivate, or so fast a rat can press a lever). Not surprisingly then, students do not know how to relate these limits to curve asymptotes. Similarly, the critical concepts of dimension and unit of measurement also are foreign to many undergraduate students of learning (and, again, we suspect, students of many other courses in psychology).

---

[6] Superstitious behavior is an interesting case. If we focus on the contingencies as the experimenter scheduled them (the viewpoint adopted in the text), then clearly we are not in the domain of causation. However, if we focus on the animal's behavior, then the difference between the effects of true causality and apparent (but false) causality may not be always obvious. But because the distinction may break down when we focus on the animal's behavior at boundary conditions (i.e., short interreinforcement intervals, brief amount of training), it does not follow that the distinction is inappropriate.

This is an unfortunate situation because the critical concepts and skills involved in thinking about variables, and drawing and interpreting functional relations, are not hard to master, at least at the level of most undergraduate psychology courses. We believe that the difficulties exist because students are rarely asked to guess how variables are functionally related, to draw a graph of the relationship, to interpret it, and to revise it. In short, the problem is, we believe, lack of practice, not the difficulty of the topic.

Here then is another great virtue of the psychology of learning: The study of learning is largely concerned with the study of functional relations, of how variables constrain each other, of how knowledge of one of them allows us to predict the other. Moreover, in contrast with research domains in which an experimenter manipulates an external variable to infer the workings of an hypothesized internal mechanism, rule, or representation, in many domains of learning the functional relations typically involve variables on the same level of analysis. For example, in the study of habituation, the investigator may relate the intensity of a stimulus, or the interstimulus interval, to the rates of response decrement and spontaneous recovery. In the study of Pavlovian conditioning, one may relate the degree of contingency between the conditional stimulus and the unconditional stimulus to the rate of acquisition and the asymptotic strength of a conditional response. In operant conditioning, reinforcement rate may be related to response rate across different reinforcement schedules. In the study of timing, the variability in the animal's behavioral estimates of a duration may be related to the magnitude of different durations. In no other area of psychology, with the exception of psychophysics, are there so many examples of empirically obtained functional relations.

It follows that the teaching of learning could be substantially enriched by exploiting the richness of its functional relations. Students could be asked to reason about a particular functional relation (e.g., how the utility of one extra dollar changes with current wealth). Intuition, elementary geometric reasoning, and some help from the instructor would allow the students to derive the typically increasing, negatively accelerated, curve. Afterwards they could also derive from the shape of the curve the predictions of utility theory concerning choice situations in which different amounts of money (or other reinforcers) can be obtained with different probabilities. We believe that, through these exercises, the resulting concepts of risk aversion, risk seeking, decreasing marginal utility, and expected value gain new meaning. More generally, we envision the day when the psychology of learning might be taught in ways similar to the current teaching of physics or chemistry. That is, having learned the fundamental principles in class, students are assigned homework problems in which they have to derive new relations, apply known concepts to new situations, and draw and interpret graphs.

Other cases could be described to illustrate how students of learning could practice with the concepts of variable and functional relation while learning the specific contents of the discipline (e.g., the matching law can be derived graphically on the basis of intuitive stability considerations; some of the predictions of the Rescorla-Wagner model can be derived on the basis of a clear understanding of the concept of steady state). But our general point is that learning textbooks and class instructors would probably achieve more if they were to inform less but form more (see also Dempster, 1993). That is, if they were (a) to substitute graphical analyses for long verbal descriptions of experiments, (b) to ask students to predict the expected functional relations before giving them the results, and (c) to better utilize opportunities to teach how to draw, read, interpret, and evaluate graphs, then they would produce some of the signs of scientific maturity. The approach we are suggesting is well summarized in the following story told by Finch (1995, p. 197): A religious teacher was once asked by a student why he did not explain the meaning of his stories more clearly, particularly because these meanings seemed so well hidden. The teacher replied, How would you feel if you asked me for an orange, and I gave you one that had already been sucked dry?

---

[7] Other common difficulties with handling graphs include dissociating the two coordinates of an *xy* graph, understanding the concept of a function as a constraint between two variables, and visualizing graphically a solution to a simple problem.

Hypothetical constructs and the unobservable. Embedded within the study of learning are two views regarding the use of intervening variables and hypothetical constructs (MacCorquodale & Meehl, 1954; see also MacCorquodale & Meehl, 1948). Traditionally, Pavlovian researchers have been accepting of constructs such as association, representation, information, surprise, and attention that deal with hypothetical processes and states of the mind or the brain. In contrast, operant researchers have tended to eschew constructs that refer to unobservable phenomena, preferring to focus on functional relations between overt behavior and operationally defined variables. Liberal or conservative uses of, and tolerances for, hypothetical constructs and intervening variables reflect a scientist's acceptance of how much guessing and speculation is allowed. These two ends of an epistemological continuum are still disputed after many years of spirited discussion (e.g., Donahoe & Palmer, 1994; Shull, 1995; Sidman, 1960; Skinner, 1950, 1990).

The student of learning could therefore be encouraged to explore the potential consequences of each extreme, the awkwardness that results when one refrains from using or inventing any intervening variables or hypothetical constructs, and the verbal noise that results when one multiplies and abuses them. Examples also abound in learning by means of which the antidote to both types of malady, a careful conceptual investigation, can usefully be taught (see, e.g., Catania's, 1975, analysis of the concept of self-reinforcement).

In summary, if epistemological issues are given the attention they deserve, from the critical importance of conceptual investigations to the mastery of the fundamental concepts of science (variables and functions), then, again, the value of the psychology of learning will be more evident.

CONCLUSIONS

To bring to light the enormous potential of the psychology of learning and end its current eclipse, instructors should be aware of how the course content can be a means toward broader, more important ends. To do so, in our view, we need to go beyond a simple definitional approach to learning, be less focused on observations, facts, procedures and results, and emphasize the ties between the issues raised in the study of learning and the issues raised in other areas of inquiry. Learning is well suited to illustrate the difficulties of understanding processes that occur on different time scales (e.g., phylogenetic and ontogenetic processes of adaptation), or the complexity of historical, strongly path-dependent systems. Learning is also well suited to explore some of the fundamental aspects of the game of science, how content problems, method, technology, and epistemology are interrelated, for example. The key concepts of variable (dimension and unit of measurement, asymptote and rate of change) and functional relation (how variables constrain each other, how they combine in algebraic structures, and the like) need to be brought to the forefront of learning textbooks and courses.

In our view, we also need to promote conceptual thinking. Science is not reducible to empirical questions, nor are all scientific problems amenable to an empirical solution. Issues dealing with the intelligibility or the sense of our concepts and their grammars, with the internal consistency of our theories, or the essential limitations of our instruments are equally important. The use of terms such as *surprise* and *attention* to understand learning may be useful, but only to the extent that both instructors and students recognize the strengths and limitations of these concepts, and that the illusion of an explanation is worse than no explanation. In short, then, we need to restore the balance between conceptual and factual investigations.

We also need to stop trivializing the subject matter by resorting to simplistic analyses of real-world examples, or by conveying to students that all is known about learning. We need to ask students to formulate questions and solutions, to think and improve upon these, and then to abandon them and search for better ones. We need to stop constantly asking students to select from a set of answers, regurgitate definitions, or reproduce procedures and results. To promote discovery—one of the rarest but also the sweetest of the scientific rewards—students need to be active. As Dethier (1962, pp. 96–97) wrote,

> Too often, ''learn'' simply means storing facts in the memory like unused furniture in a dusty attic. One of the most discouraging features about teaching is the widespread conviction among students that the process of learning is something bequeathed by the teacher, by a book, or by TV to a student. It is a startling revelation to many that learning is an active endogenous process, that learning is done by the individual by himself.

We believe that to prevent the psychology of learning from being a string of raw facts, commonsense applications, esoteric language, and (from the students' perspective) boring content, instructors should be aware that the course can be a means toward broad, important, and exciting ends. With these ends in mind we may change the miserly conception of the subject that we inadvertently helped to create, and illustrate the greatness of the subject to a new generation of undergraduates.

REFERENCES

Allport, G. W. (1961). *Pattern and growth in personality.* New York: Holt.

Bachelard, G. (1975). *La formation de l'esprit scientifique* [The formation of the scientific mind]. Paris: J. Vrin.

Baldwin, J. D., & Baldwin, J. I. (1998). *Behavior principles in everyday life* (3rd ed.). Upper Saddle River, NJ: Prentice Hall.

Barker, L. (1997). *Learning and behavior* (2nd ed.). Upper Saddle River, NJ: Prentice Hall.

Bell, M. C., & Goodie, A. S. (1997). A comparative survey of job prospects for the period 1991–1996. *APS Observer, 10,* 16–18.

Bernard, C. (1927). *An introduction to the study of experimental medicine.* New York: Macmillan. (Original work published 1865)

Blumberg, M. S., & Wasserman, E. A. (1995). Animal mind and the argument from design. *American Psychologist, 50,* 133–144.

Breland, K., & Breland, M. (1961). The misbehavior of organisms. *American Psychologist, 16,* 681–684.

Calvin, W. (1987). The brain as a Darwin machine. *Nature, 330,* 33–34.

Campbell, D. T. (1960). Blind variation and selective retention in creative thought as in other knowledge processes. *Psychological Review, 67,* 380–400.

Cantril, H. (1977). Psychology and scientific inquiry. In R. G. Colodny (Ed.), *Logic, laws, and life: Some philosophical complications* (pp. 173–183) (Vol. 6 in University of Pittsburgh Series in the Philosophy of Science). Pittsburgh: University of Pittsburgh Press.

Carnap, R. (1966). *An introduction to the philosophy of science.* New York: Dover.

Catania, A. C. (1975). The myth of self-reinforcement. *Behaviorism, 3,* 192–199.

Catania, A. C. (1998). *Learning* (4th ed.). Upper Saddle River, NJ: Prentice Hall.

Church, R. M., & Meck, W. H. (1982). Temporal generalization. *Journal of Experimental Psychology: Animal Behavior Processes, 8,* 165–186.

Cook, T. D., & Campbell, D. T. (1979). *Quasi-experimentation: Design and analysis issues in field settings.* Boston: Houghton Mifflin.

Cooper, L. D. (1991). Temporal factors in classical conditioning. *Learning and Motivation, 22,* 129–152.

Cronbach, L. J. (1957). The two disciplines of scientific psychology. *American Psychologist, 12,* 671–684.

Darwin, C. (1881). *The formation of vegetable mould through the action of worms.* London: John Murray.

Dawkins, R. (1976). *The selfish gene.* London: Oxford.

Dawkins, R. (1986). *The blind watchmaker.* New York: Norton.

Dempster, F. N. (1993). Exposing our students to less should help them learn more. *Phi Delta Kappan, 74,* 432–437.

Dethier, V. G. (1962). *To know a fly.* New York: McGraw-Hill.

Dinsmoor, J. A. (1986). Behaviorism and the education of psychologists. *Behavioral and Brain Sciences, 9,* 702.

Dinsmoor, J. A. (1989). Keller and Schoenfeld's *Principles of Psychology. The Behavior Analyst, 12,* 213–219.

Domjan, M. (1993). *The principles of learning and behavior* (3rd ed.). Pacific Grove, CA: Brooks/Cole.

Donahoe, J. W., & Palmer, D. C. (1994). *Learning and complex behavior.* Boston: Allyn & Bacon.

Drickamer, L. C., Vessey, S. H., & Meikle, D. (1996). *Animal behavior: Mechanisms, ecology, evolution* (4th ed.). Boston: Wm. C. Brown.

Eddington, A. (1939). *The philosophy of physical science.* New York: Macmillan.

Falconer, D. S., & Mackay, T. F. C. (1989). *An introduction to quantitative genetics* (4th ed.). Essex, England: Longman.

Finch, H. L. (1995). *Wittgenstein.* Rockport, MA: Element.

Gibbon, J. (1977). Scalar expectancy theory and Weber's law in animal timing. *Psychological Review, 84,* 279–335.

Glenn, S. S. (1991). Contingencies and metacontingencies: Relations among behavioral, cultural, and biological evolution. In P. A. Lamal (Ed.), *Behavioral analysis of societies and cultural practices* (pp. 39–73). New York: Hemisphere.

Grant, L., & Evans, A. (1994). *Principles of behavior analysis.* New York: Harper Collins.

Hammond, L. J. (1980). The effect of contingency upon appetitive conditioning of free operant behavior. *Journal of the Experimental Analysis of Behavior, 34,* 297–304.

Kazdin, A. E. (1982). *Single-case research designs: Methods for clinical and applied settings.* New York: Oxford University Press.

Keller, F., & Schoenfeld, W. N. (1995). *Principles of psychology.* Acton, MA: Copley Publishing Group. (Original work published 1950)

Kennedy, J. S. (1992). *The new anthropomorphism.* New York: Cambridge University Press.

Killeen, P. R. (1978). Superstition: A matter of bias, not detectability. *Science, 199,* 88–90.

Killeen, P. R. (1992). Mechanics of the animate. *Journal of the Experimental Analysis of Behavior, 57,* 429–463.

Killeen, P. R., & Fetterman, J. G. (1988). A behavioral theory of timing. *Psychological Review, 95,* 274–295.

Klein, S. B. (1996). *Learning: Principles and applications* (3rd ed.). New York: McGraw-Hill.

Lieberman, D. A. (1993). *Learning: Behavior and cognition* (2nd ed.). Pacific Grove, CA: Brooks/Cole.

MacCorquodale, K., & Meehl, P. E. (1948). On a distinction between hypothetical constructs and intervening variables. *Psychological Review, 55,* 95–107.

MacCorquodale, K., & Meehl, P. E. (1954). Edward C. Tolman. In W. K. Estes, S. Koch, K. MacCorquodale, P. E. Meehl, C. G. Mueller, W. N. Schoenfeld, & W. S. Verplank, *Modern learning theory* (pp. 117–266). New York: Appleton-Century-Crofts.

Machado, A. (1992). Behavioral variability and frequency-dependent selection. *Journal of the Experimental Analysis of Behavior, 58,* 241–263.

Machado, A. (1997). Learning the temporal dynamics of behavior. *Psychological Review, 104,* 241–265.

Martin, P., & Bateson, P. (1993). *Measuring behavior: An introductory guide* (2nd ed.). Cambridge: Cambridge University Press.

Mazur, J. E. (1998). *Learning and behavior* (4th ed.). Englewood Cliffs, NJ: Prentice Hall.

Meehl, P. (1978). Theoretical risks and tabular asterisks: Sir Karl, Sir Ronald, and the slow progress of soft psychology. *Journal of Consulting and Clinical Psychology, 46,* 806–834.

Michael, J. (1993). *Concepts and principles of behavior analysis.* Kalamazoo, MI: Association for Behavior Analysis.

Midgley, M., Lea, S. E. G., & Kirby, R. M. (1989). Algorithmic shaping and misbehavior in the acquisition of token deposit by rats. *Journal of the Experimental Analysis of Behavior, 52,* 27–40.

Morris, R. G. M. (1981). Spatial localization does not require the presence of a local cue. *Learning and Motivation, 12,* 239–260.

Olton, D. S., & Samuelson, R. J. (1976). Remembrance of places past: Spatial memory in rats. *Journal of Experimental Psychology: Animal Behavior Processes, 2,* 97–116.

Page, S., & Neuringer, A. (1985). Variability as an operant. *Journal of Experimental Psychology: Animal Behavior Processes, 11,* 429–452.

Pear, J. J., & Legris, J. A. (1987). Shaping by automated tracking of an arbitrary operant response. *Journal of the Experimental Analysis of Behavior, 47,* 241–247.

Pierce, W. D., & Epling, W. F. (1988). *Biobehaviorism: Genes, learning, and behavior* (Working paper No. 88-5). Edmonton: Center for Systems Research, University of Alberta.

Pierce, W. D., & Epling, W. F. (1995). *Behavior analysis and learning.* Englewood Cliffs, NJ: Prentice Hall.

Pinker, S. (1997, October). Against nature. *Discover, 18,* 92–95.

Plotkin, H. (1993). *Darwin machines and the nature of knowledge.* Cambridge, MA: Harvard University Press.

Popper, K. (1972). *Objective knowledge: An evolutionary approach.* Oxford: Oxford University Press.

Robinson, D. N. (1979). The history of psychology and the ends of instruction. *Teaching of Psychology, 6,* 4–6.

Romanes, G. J. (1883). *Animal intelligence.* London: Routledge & Kegan Paul.

Roughgarden, J. (1996). *Theory of population genetics and evolutionary ecology: An introduction.* Upper Saddle River, NJ: Prentice Hall.

Schwartz, B. (1982). Failure to produce response variability with reinforcement. *Journal of the Experimental Analysis of Behavior, 33,* 171–181.

Schwartz, B. (1984). *Psychology of learning and behavior.* New York: Norton.

Shull, R. L. (1995). Interpreting cognitive phenomena: Review of Donahoe and Palmer's *Learning and Complex Behavior. Journal of the Experimental Analysis of Behavior, 63,* 347–358.

Sidman, M. (1960). *Tactics of scientific research: Evaluating experimental data in psychology.* New York: Basic Books.

Silva, F. J. (in press). Beyond the "hot-and-cold" game: A demonstration of computer-controlled shaping. *Behavior Research Methods, Instruments, & Computers.*

Skinner, B. F. (1948). "Superstition" in the pigeon. *Journal of Experimental Psychology, 38,* 168–172.

Skinner, B. F. (1950). Are theories of learning necessary? *Psychological Review, 57,* 193–216.

Skinner, B. F. (1956). A case history in scientific method. *American Psychologist, 11,* 221–233.

Skinner, B. F. (1957). *Verbal behavior.* New York: Appleton-Century-Crofts.

Skinner, B. F. (1969). *Contingencies of reinforcement: A theoretical analysis.* New York: Appleton-Century-Crofts.

Skinner, B. F. (1981). Selection by consequences. *Science, 213,* 501–504.

Skinner, B. F. (1990). Can psychology be a science of mind? *American Psychologist, 42,* 780–786.

Slavin, R. E. (1997). *Educational psychology: Theory and practice.* Boston: Allyn & Bacon.

Small, W. S. (1899). An experimental study of the mental process of the rat. *American Journal of Psychology, 11,* 131–165.

Smith, K. D. (1992). On prediction and control: B. F. Skinner and the technological idea of science. *American Psychologist, 47,* 216–223.

Staddon, J. E. R. (1977). Schedule-induced behavior. In W. K. Honig & J. E. R. Staddon (Eds.), *Handbook of operant behavior* (pp. 125–152). Englewood Cliffs, NJ: Prentice Hall.

Staddon, J. E. R., & Ettinger, R. H. (1989). *Learning: An introduction to the principles of adaptive behavior.* San Diego: Harcourt Brace Jovanovich.

Staddon, J. E. R., & Simmelhag, V. (1971). The superstition experiment: A reexamination of its implications for the principles of adaptive behavior. *Psychological Review, 78,* 3–43.

Sternberg, R. (1997). Fads in psychology: What can we do. *APA Monitor, 28,* 19.

Stokes, P. D., & Balsam, P. D. (1991). Effects of reinforcing preselected approximations on the topography of the rat's bar press. *Journal of the Experimental Analysis of Behavior, 55,* 213–231.

Tarpy, R. M. (1997). *Contemporary learning theory and research.* New York: McGraw-Hill.

Timberlake, W. (1990). Natural learning in laboratory paradigms. In D. A. Dewsbury (Ed.), *Contemporary issues in comparative psychology* (pp. 31–54). Sunderland, MA: Sinauer.

Timberlake, W., & Lucas, G. A. (1985). The basis of superstitious behavior: Chance contingency, stimulus substitution, or appetitive behavior? *Journal of the Experimental Analysis of Behavior, 44,* 279–299.

Timberlake, W., & Silva, F. J. (1994). Observation of behavior, inference of function, and study of learning. *Psychonomic Bulletin & Review, 1,* 73–88.

Warner, L. H., & Warden, C. J. (1927). The development of the standardized animal maze. *Archives of Psychology, 92,* 1–35.

Whitehead, A. N. (1948). *An introduction to mathematics.* London: Oxford University Press.

Wittgenstein, L. (1958). *Philosophical investigations.* Englewood Cliffs, NJ: Prentice Hall.

Woolfolk, A. E. (1998). *Educational psychology* (7th ed.). Boston: Allyn & Bacon.

Received December 9, 1997
Final acceptance June 2, 1998